

Towards Practicable Defeasible Reasoning for ABoxes (Extended Abstract)

Jonas Haldimann^{1,2}, Magdalena Ortiz¹ and Mantas Šimkus¹

¹Institute of Logic and Computation, TU Wien, 1040 Vienna, Austria

²University of Cape Town and CAIR, Cape Town, South Africa

Abstract

Reasoning with rules that allow for exceptions has been a longstanding challenge in knowledge representation. The KLM paradigm has been successful for defeasible reasoning in propositional logics, but its application to Description Logics (DLs) has been challenging. Many approaches to terminological reasoning with defeasible inclusions have been proposed, but reasoning about ABoxes is still largely unexplored. In this paper, we consider defeasible inclusions in the expressive DL \mathcal{ALCT} with closed predicates, but restrict the inclusions in a way that circumvents some of the challenges faced by related approaches. We also consider the data complexity of defeasible reasoning, which, to our knowledge, had not yet been analysed. Unfortunately, our approach is hard for the second level of the polynomial hierarchy, but we identify a restricted fragment that enables tractable reasoning.


Keywords


defeasible reasoning, rational closure, ABoxes, complexity of reasoning


1. Introduction

Defeasible reasoning, that is, reasoning with rules that allow for exceptions, has been a longstanding challenge in knowledge representation. The KLM paradigm for defeasible reasoning [1] has been successful in propositional logics [2, 3, 4, 5], providing a principled approach that is not as computationally expensive as other NMR formalisms. Extending this approach to DLs is clearly appealing, and it has received significant attention in the literature [6, 7, 8, 9, 10, 11, 12]. But the focus of these works has been almost exclusively on inferring defeasible inclusions from TBoxes. Data-centric reasoning services, such as defeasible instance checking in the presence of ABoxes, have been largely overlooked and, to our knowledge, the data complexity of defeasible inferences about ABox objects remained unexplored.

In our recent paper [13], we consider the expressive DL \mathcal{ALCT} with *closed predicates*, which already allows some simple non-monotonic reasoning and which generalizes nominals [14, 15]; this is one of the most expressive DLs that can be decided in ExpTime . We add defeasible inclusions to \mathcal{ALCT} knowledge bases and give them an *exceptionality based* semantics in the style of *Rational Closure* [16] and the equivalent *system Z* [3]. The classical part of the knowledge base (KB) remains unrestricted, providing the full power of the DL to draw conclusions about both named and unnamed objects, but the defeasible inclusions are syntactically restricted in a way that they can draw inferences only about the ABox individuals. The result is a simple formalism that allows to draw defeasible inferences about specific objects, and where the combined complexity of these inferences is not higher than for classical reasoning in the underlying DL. Unfortunately, this is not the case for data complexity: we show that credulous and sceptical entailment are both intractable. Simply restricting the DL is not enough to regain tractability, since changing the assignment of defeasible concepts to an individual can affect the defeasible assignment of its neighbours, which rules out the existence of a unique preferred model. Nevertheless, we identify a restricted fragment that allows for defeasible instance checking in a time that is polynomial in the size of the ABox.

 DL 2025: 38th International Workshop on Description Logics, September 3–6, 2025, Opole, Poland

 jonas@haldimann.de (J. Haldimann); magdalena.ortiz@tuwien.ac.at (M. Ortiz); mantas.simkus@tuwien.ac.at (M. Šimkus)

 0000-0002-0877-7063 (J. Haldimann); 0000-0002-2344-9658 (M. Ortiz); 0000-0003-0632-0294 (M. Šimkus)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Existing defeasible semantics for DLs based on rational closure neglect all defeasible information for the unnamed objects implied by the existential axioms, sometimes leading to counterintuitive inferences. To our knowledge, the only adequate solution so far is limited to the very inexpressive \mathcal{EL} . Our approach fully circumvents the issue of typicality of anonymous objects: since defeasible inferences do not allow inferring new facts about unnamed objects, there is no need to decide how to apply the defeasible inclusions to them. While we do not solve the problem, we do avoid its nontransparent and sometimes counterintuitive aspects. We believe this is a reasonable compromise. Unnamed objects in an interpretation intuitively describe structures that should be as general as possible, and we have not seen many realistic examples that really call for defeasible inferences over anonymous objects.

2. The Formalism

In this section, we formally define our DL knowledge bases with *defeasible inclusions* which are concept inclusions that hold for typical, but not necessarily all elements of the domain. Unlike a ‘strict’ concept inclusion $C \sqsubseteq D$, a defeasible inclusion $C \sqsubset D$ may have exceptions, i.e., in a model \mathcal{I} of $C \sqsubset D$ some elements of $C^{\mathcal{I}}$ may not be in $D^{\mathcal{I}}$; these are considered *exceptional*.

Our key design choice is to make sure that defeasible inclusions apply only to objects that are explicitly named in the knowledge base, and therefore anonymous objects whose existence may be implied by concept inclusions can be treated as non-exceptional. This will be ensured by means of *rooted concepts*, which are concept expressions that are ‘guarded’ by closed predicates. For closed predicates C and r in $N_C^c \subseteq N_C$ and $N_R^c \subseteq N_R$, respectively, the term $\mathcal{I} \models^c \mathcal{A}$ means that $C^{\mathcal{I}}, r^{\mathcal{I}}$ only contain an individual if explicitly mentioned. In our formalisms, defeasible inclusions are restricted to have rooted concepts in the antecedents.

Definition 1. A concept C is rooted if one the following is satisfied: 1. $C \in N_C^c$, 2. C of the form $C_1 \sqcap C_2$ and at least one of C_1, C_2 is rooted, 3. C of the form $C_1 \sqcup C_2$ and both C_1, C_2 are rooted, or 4. C of the form $\exists r.D$ with $r \in N_R^c$ or $r^- \in N_R^c$.

A defeasible concept inclusion (DCI) is an expression $C \sqsubset D$, where C, D are concepts and C is rooted. A knowledge base (KB) is a tuple $\mathcal{K} = (\mathcal{T}, \mathcal{D}, \mathcal{A})$, where \mathcal{T} is a TBox, \mathcal{D} is a set of defeasible inclusions, and \mathcal{A} is an ABox.

Example 1 (adapted from [6]). Consider the knowledge that red blood cells (RBC) usually have a nucleus, but mammalian red blood cells (MRBC) typically do not have a nucleus; and of course mammalian red blood cells are a subclass of red blood cells. c_1, c_2, c_3 are red blood cells and c_3 is additionally a mammalian red blood cell. Furthermore, we assume the knowledge about RBC and MRBC to be complete for c_1, c_2, c_3 . This situation can be described by the knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{D}, \mathcal{A})$ with $\mathcal{A} = \{RBC(c_1), RBC(c_2), RBC(c_3), MRBC(c_3)\}$, $\mathcal{T} = \{MRBC \sqsubseteq RBC, \exists hasNucleus.\top \sqsubseteq EN\}$, and $\mathcal{D} = \{RBC \sqsubset \exists hasNucleus.\top, MRBC \sqsubset \neg \exists hasNucleus.\top\}$ where $MRBC$ and RBC are closed predicates.

We will now define the semantics of a KB $\mathcal{K} = (\mathcal{T}, \mathcal{D}, \mathcal{A})$ as the interpretations \mathcal{I} where $\mathcal{I} \models \mathcal{T}$, $\mathcal{I} \models^c \mathcal{A}$, and \mathcal{I} complies with the defeasible inclusions in \mathcal{D} as much as possible. To properly define the latter, following the definition of system Z, the first step is to define the notion of *tolerance*, which generalizes a similar concept for propositional rules [3].

Definition 2. We write $\mathcal{I}, e \models C \sqsubset D$, if $e \in (\neg C \sqcup D)^{\mathcal{I}}$. For a set \mathcal{D} of defeasible inclusions, we write $\mathcal{I}, e \models \mathcal{D}$, if $\mathcal{I}, e \models C \sqsubset D$ for all $C \sqsubset D \in \mathcal{D}$. A defeasible inclusion $C \sqsubset D$ is tolerated by a set \mathcal{D} of defeasible inclusions and a TBox \mathcal{T} , if there is an interpretation \mathcal{I} and an object $e \in \Delta^{\mathcal{I}}$ such that $\mathcal{I} \models \mathcal{T}$, $e \in (C \sqcap D)^{\mathcal{I}}$ and $\mathcal{I}, e \models \mathcal{D}$.

For a set \mathcal{D} of defeasible inclusions and a TBox \mathcal{T} construct the following sequence:

- (i) Let \mathcal{D}^0 contain all $C \sqsubset D \in \mathcal{D}$ such that $C \sqsubset D$ is tolerated by \mathcal{D} and \mathcal{T} .
- (ii) For all $\ell > 0$, let \mathcal{D}^ℓ contain all $C \sqsubset D \in \mathcal{D}'_\ell$ such that $C \sqsubset D$ is tolerated by \mathcal{D}'_ℓ and \mathcal{T} , where $\mathcal{D}'_\ell = \mathcal{D} \setminus (\mathcal{D}^0 \cup \dots \cup \mathcal{D}^{\ell-1})$.

Let k be the smallest integer such that $\mathcal{D}^{k+1} = \emptyset$. Then $(\mathcal{D}^0, \dots, \mathcal{D}^k, \mathcal{D}^\infty)$ with $\mathcal{D}^\infty = \mathcal{D}'_{k+1}$ is called the tolerance partition of \mathcal{D} (w.r.t. \mathcal{T}).

Let \mathcal{I} be an interpretation such that $\mathcal{I} \models \mathcal{D}^\infty \cup \mathcal{T}$ and assume $e \in \Delta^\mathcal{I}$. We let $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e)$ be defined as follows. If $\mathcal{I}, e \models \mathcal{D}$, then $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e) = 0$. Otherwise, $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e)$ is the biggest $i \in \{1, \dots, k+1\}$ such that $\mathcal{I}, e \not\models \mathcal{D}^{i-1}$.

Intuitively, $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e)$ tells us to what extent the defeasible inclusions are satisfied at e in \mathcal{I} . If $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e) = 0$, then e is non-exceptional and satisfies all inclusions in \mathcal{D} . If $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e) = k+1$, then e is highly exceptional: it violates some inclusion in \mathcal{D}^k , which stores the most specific defeasible inclusions of \mathcal{D} . Note that the rootedness condition guarantees that unnamed objects have $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, e) = 0$.

Let $\Delta_{\mathcal{A}}$ be the set of individuals occurring explicitly in an ABox \mathcal{A} . We can now compare the extent to which interpretations satisfy defeasible inclusions:

Definition 3. Assume a KB $\mathcal{K} = (\mathcal{T}, \mathcal{D}, \mathcal{A})$. Let $(\mathcal{D}^0, \dots, \mathcal{D}^k, \mathcal{D}^\infty)$ be the tolerance partition of \mathcal{D} w.r.t. \mathcal{T} . An interpretation \mathcal{I} is called \mathcal{K} -admissible, if $\mathcal{I} \models \mathcal{T}$, $\mathcal{I} \models^c \mathcal{A}$, and $\mathcal{I} \models \mathcal{D}^\infty$. Assume a pair \mathcal{I}, \mathcal{J} of \mathcal{K} -admissible interpretations. We write $\mathcal{I} \prec_{\mathcal{K}} \mathcal{J}$, if the following holds:

- $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, a) \leq \text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{J}, a)$ for all individuals $a \in \Delta_{\mathcal{A}}$, and
- $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, a) < \text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{J}, a)$ for some individual $a \in \Delta_{\mathcal{A}}$.

A \mathcal{K} -admissible interpretation \mathcal{J} is called a minimal model of \mathcal{K} , if there exists no \mathcal{K} -admissible interpretation \mathcal{I} such that $\mathcal{I} \prec_{\mathcal{K}} \mathcal{J}$.

Example 2. Consider \mathcal{K} from Example 1. The tolerance partition of \mathcal{D} is $(\mathcal{D}^0, \mathcal{D}^1)$ with $\mathcal{D}^0 = \{RBC \sqsubseteq \exists \text{hasNucleus}.\top\}$ and $\mathcal{D}^1 = \{MRBC \sqsubseteq \neg \exists \text{hasNucleus}.\top\}$.

Let \mathcal{I} be an interpretation with $\mathcal{I} \models \mathcal{A} \cup \mathcal{T}$ and $\mathcal{I} \models (\exists \text{hasNucleus}.\top)(c_1)$, $\mathcal{I} \models (\neg \exists \text{hasNucleus}.\top)(c_2)$, $\mathcal{I} \models (\exists \text{hasNucleus}.\top)(c_3)$. We have $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, c_1) = 0$, since $\mathcal{I}, c_1 \models \mathcal{D}$. We have $\text{rank}_{\mathcal{D}, \mathcal{T}}(\mathcal{I}, c_2) = 1$, because \mathcal{I} violates an inclusion in \mathcal{D}^0 for c_2 .

Observe that even KBs that have a canonical model may have more than one minimal model. Therefore, we consider sceptical and credulous entailment of assertions in the minimal models. *Sceptical entailment* over the minimal models accepts only conclusions that hold in all minimal models, while *credulous entailment* accepts conclusions that hold in at least one minimal model.

For this formalism, we show the following complexity results.

Theorem 1. *Credulous and sceptical entailment of assertions is EXPTIME-complete in combined complexity. In data complexity, credulous and sceptical entailment of assertions is Σ_2^P -complete and Π_2^P -complete, respectively.*

This complexity is too high for applications involving large ABoxes, hence we identify a restricted fragment of our formalism that has tractable data complexity.

Definition 4 (Local KBs). A complex concept C is closed if all concept and role names occurring in it are closed. An ordinary inclusion $C \sqsubseteq D$ or a defeasible inclusion $C \sqsubseteq \mathcal{D} D$ is called local, if every quantified concept of the form $\exists r.E$ or $\forall r.F$ occurring in C or D is closed. A knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{D}, \mathcal{A})$ is local if every defeasible inclusion in \mathcal{D} and every inclusion in \mathcal{T} is local.

The intuition of local inclusions is that they can only describe an object and its immediate surroundings. While they can use quantifiers on closed roles and predicates, which can be evaluated by looking up the assertions in the ABox, they cannot be affected by the assignment of open concepts in neighbouring nodes. This will allow us to answer queries about an object without the need to consider the assignments of open concepts for all objects in \mathcal{A} .

Theorem 2. *Let \mathcal{K} be a local KB and α be an assertion. Checking whether \mathcal{K} sceptically (or credulously, resp.) entails $C(a)$ is polynomial in data complexity and in P^{NP} in combined complexity.*

3. Related Work and Conclusions

We have presented a defeasible reasoning framework over ABoxes based on rational closure, identifying a tractable fragment where queries can be efficiently answered.

Despite the many works that extend DLs with defeasible reasoning based on rational closure, reasoning about ABoxes is still lacking. While rational closure extends well to DLs with the disjoint model union property [10, 17], it struggles with DLs that include individuals and closed predicates. Stable rational closure has been proposed for more expressive logics like *SR_{OLQ}* [6], and although our approach appears compatible in specific cases, a thorough comparison remains pending. Unlike most related methods, our framework avoids the quantification neglect problem by design, since it applies only to named individuals in the ABox. Prior solutions to this problem exist only for lightweight logics such as \mathcal{EL}_\perp [7].

While our method sidesteps the issue of quantification neglect, it does suffer from other limitations of rational closure, like inheritance blocking. Future work includes exploring more general solutions to quantification neglect and how alternative semantics could be adapted for tractable reasoning, potentially inspired by approaches like \mathcal{DL}^N [18, 19].

Acknowledgments

This work was partially supported by the Austrian Science Fund (FWF) projects PIN8884924, P30873 and 10.55776/COE12.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, *Artif. Intell.* 44 (1990) 167–207.
- [2] E. W. Adams, *The Logic of Conditionals: An Application of Probability to Deductive Logic*, Synthese Library, Springer Science+Business Media, Dordrecht, 1975.
- [3] J. Pearl, System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: R. Parikh (Ed.), *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, Pacific Grove, CA, USA, March 1990, Morgan Kaufmann, 1990, pp. 121–135.
- [4] D. Lehmann, Another perspective on default reasoning, *Ann. Math. Artif. Intell.* 15 (1995) 61–82. doi:10.1007/BF01535841.
- [5] C. Komo, C. Beierle, Nonmonotonic reasoning from conditional knowledge bases with system W, *Ann. Math. Artif. Intell.* 90 (2022) 107–144. doi:10.1007/s10472-021-09777-9.
- [6] P. A. Bonatti, Rational closure for all description logics, *Artif. Intell.* 274 (2019) 197–223. doi:10.1016/J.ARTINT.2019.04.001.
- [7] M. Pensel, A. Turhan, Reasoning in the defeasible description logic ϵ - computing standard inferences under rational and relevant semantics, *Int. J. Approx. Reason.* 103 (2018) 28–70. doi:10.1016/J.IJAR.2018.08.005.
- [8] G. Casini, T. A. Meyer, K. Moodley, U. Sattler, I. Varzinczak, Introducing defeasibility into OWL ontologies, in: M. Arenas, Ó. Corcho, E. Simperl, M. Strohmaier, M. d’Aquin, K. Srinivas, P. Groth, M. Dumontier, J. Heflin, K. Thirunarayan, S. Staab (Eds.), *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference*, Bethlehem, PA, USA, October 11–15, 2015, *Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 409–426. doi:10.1007/978-3-319-25010-6_27.

- [9] G. Casini, U. Straccia, Defeasible inheritance-based description logics, *J. Artif. Intell. Res.* 48 (2013) 415–473. doi:10.1613/JAIR.4062.
- [10] K. Britz, G. Casini, T. Meyer, K. Moodley, U. Sattler, I. Varzinczak, Principles of klm-style defeasible description logics, *ACM Trans. Comput. Log.* 22 (2021) 1:1–1:46. doi:10.1145/3420258.
- [11] G. Casini, U. Straccia, A rational entailment for expressive description logics via description logic programs, in: E. Jembere, A. J. Gerber, S. Viriri, A. W. Pillay (Eds.), *Artificial Intelligence Research - Second Southern African Conference, SACAIR 2021, Durban, South Africa, December 6-10, 2021, Proceedings*, volume 1551 of *Communications in Computer and Information Science*, Springer, 2021, pp. 177–191. doi:10.1007/978-3-030-95070-5_12.
- [12] L. Giordano, V. Gliozzi, Reasoning about exceptions in ontologies: from the lexicographic closure to the skeptical closure, *Fundam. Informaticae* 176 (2020) 235–269. doi:10.3233/FI-2020-1973.
- [13] J. Haldimann, M. Ortiz, M. Šimkus, Towards practicable defeasible reasoning for ABoxes, in: *Logics in Artificial Intelligence - 19th European Conference, JELIA 2025*, 2025.
- [14] E. Franconi, Y. A. Ibáñez-García, Í. Seylan, Query answering with DBoxes is hard, *Electronic Notes in Theoretical Computer Science* 278 (2011) 71–84. doi:<https://doi.org/10.1016/j.entcs.2011.10.007>, proceedings of the 7th Workshop on Methods for Modalities (M4M'2011) and the 4th Workshop on Logical Aspects of Multi-Agent Systems (LAMAS'2011).
- [15] C. Lutz, I. Seylan, F. Wolter, The data complexity of ontology-mediated queries with closed predicates, *Logical Methods in Computer Science* Volume 15, Issue 3 (2019). doi:10.23638/LMCS-15(3:23)2019.
- [16] D. Lehmann, What does a conditional knowledge base entail?, in: R. J. Brachman, H. J. Levesque, R. Reiter (Eds.), *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*. Toronto, Canada, May 15-18 1989, Morgan Kaufmann, 1989, pp. 212–222.
- [17] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, Semantic characterization of rational closure: From propositional logic to description logics, *Artif. Intell.* 226 (2015) 1–33. doi:10.1016/j.artint.2015.05.001.
- [18] P. A. Bonatti, M. Faella, I. M. Petrova, L. Sauro, A new semantics for overriding in description logics, *Artif. Intell.* 222 (2015) 1–48. doi:10.1016/J.ARTINT.2014.12.010.
- [19] P. A. Bonatti, L. Sauro, On the logical properties of the nonmonotonic description logic dl^n , *Artif. Intell.* 248 (2017) 85–111. doi:10.1016/J.ARTINT.2017.04.001.